

Synchronization of finite-state pulse-coupled oscillators

Hanbaek Lyu¹

Department of Mathematics, The Ohio State University, Columbus, OH 43210 (yu.1242@osu.edu)

Abstract

We propose a novel generalized cellular automaton(GCA) model for discrete-time pulse-coupled oscillators and study the emergence of synchrony. Given a finite simple graph and an integer $n \geq 3$, each vertex is an identical oscillator of period n with the following weak coupling along the edges: each oscillator inhibits its phase update if it has at least one neighboring oscillator at a particular "blinking" state and if its state is ahead of this blinking state. We obtain conditions on initial configurations and on network topologies for which states of all vertices eventually synchronize. We show that our GCA model synchronizes arbitrary initial configurations on paths, trees, and with random perturbation, any connected graph. In particular, our main result is the following local-global principle for tree networks: for $n \in \{3, 4, 5, 6\}$, any n -periodic network on a tree synchronizes arbitrary initial configuration if and only if the maximum degree of the tree is less than the period n .

Keywords:

Synchronization, pulse-coupled oscillators, generalized cellular automata, digital clock synchronization, self-stabilization, path, tree, absorbing chain

1. Introduction

The emergence of collective behavior from locally interacting many-agent systems is a pervasive phenomenon in nature and has raised strong scientific interests [28]. In particular, synchronization in systems of pulse-coupled oscillators(PCOs) is a fundamental issue in physics, biology, and engineering. Examples from nature include synchronization of fireflies [5], neurons in the brain [31], and the circadian pacemaker cells [11]). These are all examples of complex systems, which achieve desired system behavior by an aggregation of decentralized interactions between individual agents. This "bottom-up" approach became popular as a paradigm for studying complex systems after the celebrated Boids model by Reynolds [26], which beautifully demonstrated the flocking of birds or fishes in such a framework. Nowadays the technique of decentralized control finds its use in cooperative control of networked dynamical systems, from robotic vehicle networks to electric power networks to synthetic biological networks ([30], [2], [3], [6], [20], [22]).

A discrete-time deterministic dynamical system on a network of finite-state machines with a locally defined homogeneous transition map is called a generalized cellular automaton(GCA), commonly known as a cellular automaton(CA) when the network topology is taken to be a lattice [35]. GCAs can exhibit striking spatio-temporal patterns in spite of their simplicity [19], and are gaining growing interest as an alternative paradigm for modeling complex

systems [7], [16]. Some extensively studied models include lattice gas automaton for simulating fluid flows [34], Greenburg-Hastings model for excitable medium [14], and Griffeath's cyclic cellular automaton, which shows clustering on \mathbb{Z} and autowave behavior on \mathbb{Z}^2 [13]. To the author's knowledge, however, there has not been a direct attempt to use GCAs to study the synchronization of PCOs.

We can roughly classify the literature on coupled non-linear oscillators by the following four parameters: time, space, coupling, and dynamics. First, most classical studies on this subject use ODE models, which have continuous-time, continuous-space, deterministic dynamics, and phase-coupling. A pioneering work was done by Winfree [33], and then Kuramoto model has become a paradigm in this area [1], [27]. In case of biological oscillators, one usually assumes the mutual coupling is episodic and pulselike ([12], [15], [24], [32]). Peskin [25] studied a system of pulse-coupled oscillators(PCOs) to model cardiac pacemaker cells, and later Mirollo and Strogatz [21] generalized his model and showed that synchronization is guaranteed for almost all initial configurations when the oscillators are all-to-all connected. A recent work of Nishimura and Friedman [23] further generalizes their model, and derives a condition on the initial configuration to guarantee synchronization for arbitrary connected topology. To obtain synchronization for both arbitrary connected topology and initial configuration, Klingmayr et. al [18] studied a discrete-time continuous-space system of PCOs with stochastic dynamics. DeVille and Peskin [8] studied all-to-all networks of a discrete and stochastic version of Peskin's excitatory model in [25].

In computer science, a challenge in discrete deterministic coupled oscillators is known as the *digital clock synchronization problem*, which is to devise a protocol to achieve synchronization on a network of digital clocks that are synchronously updated. Two properties are highly desirable for such protocols: 1) it uses constant number of states on each local clock to ensure scalability, and 2) it synchronizes arbitrary configuration for initial synchronization and fault-tolerance. The second property is called self-stabilization, first proposed by Dijkstra [9], which is equivalent to requiring that the set of desirable system states for a protocol is a global attractor in the corresponding discrete dynamical system. Such protocols are readily applicable in wireless sensor networks, for example [29]. Nevertheless, it is well-known that there is no such protocol with both properties 1) and 2) that works on arbitrary connected networks [10]. While there does exist a protocol with property 2) for arbitrary network topology but with a bounded number of states that depends on the network [4], the more relevant work is done by Herman and Ghosh [17]; they presented a 3-state protocol on trees, which can be regarded as a 3-state GCA model for *phase*-coupled oscillators.

In this paper, we propose GCA models for a discrete system of PCOs and study their network behavior. We call our models the *firefly networks*, due to our initial motivation to understand the emergence of synchronous blinking of fireflies. These are discretized versions of previously studied models mentioned earlier (in particular, the continuous model of Nishimura and Friedman [23]) and enjoy some of the similar behavior(in particular, Lemma 2.2). Moreover, our model can be regarded as a protocol for digital clock-synchronization that uses constant number of states. Our main results tell us that for some classes of network topologies synchrony is guaranteed to emerge, but there are also examples of connected networks where synchrony may fail to emerge. These results contrast with the models that incorporate a certain type of stochasticity mentioned earlier, for which emergence of synchrony with probability 1 is derived for all connected network of finitely many oscillators [18]. We also obtain universal synchrony for a randomized version of our model as a corollary of our deterministic results.

In the rest of the introduction, we give a definition of our GCA model together with some illustrating examples and our main results. Some of the results are derived for the firefly system defined here, and some for more general types of GCA.

Definition 1. Let $G = (V, E)$ be a finite simple graph and fix $n \geq 3$. Let $\mathbb{Z}_n = \mathbb{Z}/n\mathbb{Z}$ with linear ordering $0 < 1 < 2 < \dots < n-1$. An n -configuration is a map $X : V \rightarrow \mathbb{Z}_n$. Let $b(n) = \lfloor \frac{n-1}{2} \rfloor$ be the blinking state. The time evolution of a given initial configuration is given by the firefly transition map $\tau_G : X \mapsto X'$, defined as follows :

$$X'(v) = \begin{cases} X(v) & \text{if } X(v) > b(n) \text{ and } v \text{ is adjacent to some vertex of state } b(n) \\ X(v) + 1 & \text{otherwise} \end{cases}. \quad (1)$$

The discrete-time dynamical system on G generated by the iteration of this transition map is called the n -periodic firefly network on G . The firefly network on G starting from an initial configuration X_0 is denoted (G, τ_G, X_0) or (G, X_0) in short, and the configuration after $t \in \mathbb{N}$ iterations of the transition map will be denoted X_t . We call the unit of time "second". The sequence $(X_t)_{t=0}^\infty$ will be called the orbit of X_0 on G . We say (G, X_0) synchronizes or X_0 synchronizes if there is $N \in \mathbb{N}$ such that X_t is a constant function for all $t \geq N$. We say G is n -synchronizing if every n -configuration on G synchronizes.

In words, our transition rule can be interpreted as follows: in a network of n -state identical oscillators, each oscillator updates from state s to $s+1 \pmod n$ unless it senses a blinking state and notices that its phase is ahead of the blinking neighbor, in which case it waits for 1 second without update. In this case, we say that the vertex is *pulled* by its blinking neighbor.

Our network model is a discrete-time dynamical system with finite configuration space. Hence, every trajectory must converge to a limit cycle. Limit cycles can be either a synchronous or asynchronous periodic orbit, as illustrated in the examples of 6-periodic networks in Figures 1 and 2. Note that $b(6) = 2$ is the blinking state in this case, so every vertex of state 3, 4, or 5 with a state 2 neighbor stops evolving for 1 second and all the other vertices evolves to the next state.




Figure 1: Two examples of synchronizing 6-periodic networks are shown in (a) and (b). In (c) and (d), the last configurations are symmetric to the initial ones, so the networks do not synchronize.



Figure 2: More examples of non-synchronizing 6-periodic firefly networks.

As illustrated in the examples in Figure 1 and 2, whether a network (G, X_0) synchronizes depends on the structure of G , initial configuration X_0 , and also on the period n as we will see at the end of Section 6. We find conditions on G , n , and X_0 to guarantee synchronization. In particular, we find conditions on G and n such that our GCA model for synchronization is self-stabilizing, i.e., it synchronizes arbitrary n -configurations on G . It appears that paths are inherent to our model, in the sense of the following theorem:

Theorem 2. *Every path P is n -synchronizing for all $n \geq 3$. Furthermore, the maximum synchronization time, i.e., the largest possible number of steps to synchronize an initial configuration is linear in the size of P . More precisely, let P be a path of m vertices, and let $T_P(m)$ be the maximum synchronization time. Then we have the following linear bound*

$$n \left(\frac{n}{2} - 1 + m \right) \leq T_P(m) \leq (m-1) \left(\frac{n^2}{2} + 2n - 2 \right). \quad (2)$$

As one tries to expand the class of n -synchronizing graphs, the class of finite trees would be a natural choice that includes finite paths. However, as illustrated by example (c) in Figure 1, there exists a tree with a bad 6-configuration, which never synchronizes. In this counterexample, the obvious obstacle to obtain synchrony on trees is that a vertex with many neighbors could stop updating when it is constantly pulled by its neighbors. Namely, let v be a vertex of a finite tree T with degree $\geq n$, and let T_1, \dots, T_m be the connected components of $T - v$, the graph with v and its edges are removed from T . Note that $m \geq n$. Assign state $i \pmod n$ to every vertex of T_i , and assign any state $> n/2$ to vertex v . Then v never blinks and each component T_i never get pulled by v , which is essentially the counterexample in Figure 1 (c). Hence if every n -configuration on T synchronizes, it is necessary that T has maximum degree $< n$.

A non-trivial result is the converse, namely, the "if" part of the following theorem

Theorem 3. *Let T be a tree and let $n \in \{3, 4, 5, 6\}$. Then T is n -synchronizing if and only if the maximum degree of T is strictly less than n .*

This result gives a necessary and sufficient local condition on network topology that guarantees global synchrony. Observe that the key feature of the counterexample Figure ?? (c) is that some vertex could stop blinking eventually due to constant pullings from its neighbors, and in turn those neighbors stop being pulled by this vertex. This isolates the connected components of $T - v$ due to the tree structure. Indeed, the following theorem shows that for certain periods, this is a crucial factor to guarantee synchrony on trees.

Theorem 4. *Let T be a tree and let X_0 be a n -configuration for some $n \in \{3, 4, 5, 6\}$. Then (T, X_0) synchronizes if and only if every vertex of T blinks infinitely often in the dynamic.*

An easy way to achieve such blinking property is to make the maximum degree of the underlying graph to be less than the period n , as in the following lemma.

Lemma 5. *Let $G = (V, E)$ be a graph and let u be a vertex. Suppose $\deg_G(u) < n$. Let X_0 be any n -configuration on G . Then u blinks infinitely often in the dynamic (G, X_0) .*

Now Theorem 3 follows easily from Theorem 4 and Lemma 5; for if $n \in \{3, 4, 5, 6\}$ and T is a tree with maximum degree $< n$, then for any n -configuration X_0 every vertex blinks infinitely often in the dynamic, and hence the configuration synchronizes by the theorem.

Our proof of Theorem 4 when $n = 6$ is long and technical, so it will be omitted in this paper. It should be mentioned that the local-global principle of Theorem 2 fails to hold for $n = 7$. Consider a star of degree 4 with following 7-periodic initial configuration in Figure 3. Every vertex blinks infinitely often in the network but it does not synchronizes. Verifying




Figure 3: A counterexample for the blinking tree theorem for $n = 7$. Diagrams in the right represents the dynamic on this counterexample in an alternative representation called relative circular representation, which will be introduced in Definition 2.3. The square and dots here denote the relative state of the center and the leaves, respectively.

Theorem 4 for $n \geq 8$ is still open.

The above results tell us that for some classes of network topologies synchrony is guaranteed to emerge, possibly depending on the period n , whereas there are also some connected networks where synchrony may fail to emerge. While it is interesting but difficult to characterize the class of synchronizing graphs, if we introduce a certain type of stochasticity in our deterministic model, we can easily show that the network on any connected graph synchronizes with probability 1. The intuition behind this randomization is that by randomness we can disregard subtle influence of the underlying topology on the global dynamic and break the symmetry of non-synchronizing orbits. The following result illustrates this observation; assuming stochastic reception of signals in our deterministic model, we quickly obtain the following result as an application of Theorem 2.

Theorem 6. *Let G be any finite connected graph and fix $n \geq 3$. Let X_0 be any n -configuration on G . Suppose that each edge e in G is present at each instant independently with a fixed probability $p_e \in (0, 1)$. Then the dynamic (G, X_0) synchronizes with probability 1.*

In section 2 we discuss a more general version of the GCA of Definition 1, give a geometric representation of the dynamic, and discuss a fundamental observation on stable manifolds on which synchronization is guaranteed for arbitrary connected topology. And then we mention a characteristic property of our model as an inhibitory system, which enables inductive arguments in the later sections. In section 3, we establish a key lemma (Lemma 3.4) in this paper and prove Theorem 2. In the following section, Section 4, we discuss transient and recurrent local configurations based on the concept of Poincaré return map, and prove Theorem 3 for $n \in \{3, 4, 5\}$. In section 5, we randomize our deterministic model and show that the resulting Markov chain is in fact an absorbing chain with synchrony being the unique absorbing

state. Hence such randomized versions of our model synchronizes with probability 1 on any connected graphs.

In the following discussions, we assume every graph is finite, simple, and connected, unless otherwise mentioned. If S, H are vertex-disjoint subgraphs of G , then $S + H$ is defined by the subgraph obtained from $S \cup H$ by adding all the edges in G between S and H . On the other hand, $S - H := S - V(H)$ denotes the subgraph of G obtained from S by deleting all vertices of H and edges incident to them. If $v \in V(H)$ and H is a subgraph of G , then $N_H(v)$ denotes the set of all neighbors of v in H and $\deg_H(v) := |N_H(v)|$.

2. Generalities, relative circular representation, and the width lemma

We begin our discussion with a general consideration on GCA models for finite-state coupled oscillators. Let $G = (V, E)$ be a graph and fix $n \geq 3$, and let $S(G, n) = \mathbb{Z}_n^{|V|}$ be the set of all n -configurations on G . A transition map τ_G on $S(G, n)$ then defines deterministic dynamics on G . Let τ_G be the firefly transition map given in Definition 1. It has the following natural properties that would be required for any GCA model for coupled oscillators:

- (i) (*isolation*) If $v \in V$ is isolated in G , then $\tau_G(X)(v) = X(v) + 1$;
- (ii) (*local dependence*) If $v \in V$ and $e \in E$ is any edge that is not incident to v , then $\tau_G(X)(v) = \tau_{G-e}(X)(v)$.

Moreover, the coupling given by τ_G is pulse-coupling by which we mean the following property:

- (iii) (*pulse-coupling*) There exists a unique state $b(n) \in \mathbb{Z}_n$, called the *blinking state*, such that for each $v \in V$, we have $\tau_G(X)(v) = X(v) + 1$ if no neighbor of v has state $b(n)$.

In words, an oscillator evolves to the next state if there is no blinking neighbor, as a non-blinking firefly will not affect its neighboring fireflies or a neuron will not be affected by its non-spiking neighboring neurons. A *discrete system of n -periodic PCO (pulse-coupled oscillator)s* is an assignment $\tau : G \mapsto \tau_G$ for each finite connected graph G to a transition map τ_G on the set of n -configurations on G satisfying conditions (i)-(iii). Hence the n -periodic firefly network defined in Definition 1 is a discrete system of PCOs with $b(n) = \lfloor \frac{n-1}{2} \rfloor$ being the blinking state.

Note that since the state space \mathbb{Z}_n^V is finite and the dynamic is deterministic, for any system of oscillators on a graph $G = (V, E)$ with any initial configuration, the trajectory must converge to a periodic orbit. The simplest possible and desired orbit is synchrony, where all oscillators have identical states. We wish to achieve global synchrony by an aggregation of local efforts to obtain mutual synchrony. Hence it is natural to require that the system τ synchronizes two coupled oscillators regardless of initial configuration. Our model indeed has this property with local monotonicity in the sense of property (iv) below:

Definition 2.1. Let $G = (V, E)$ be a graph, and let $X : V \rightarrow \mathbb{Z}_n$ be an n -configuration. Let u, v be two vertices in G . The clockwise displacement of v from u in configuration X is defined by

$$\delta_X(u, v) := X(v) - X(u) \pmod{n}.$$

If X_0 is an initial configuration, we write $\delta_t(u, v) := \delta_{X_t}(u, v)$ for all $t \geq 0$. We say v is clockwise to u and u is counterclockwise to v at t if $\delta_t(u, v) < n/2$, and u is opposite to v if $\delta_t(u, v) = n/2$,

which can happen only if n is even. Suppose u and v are adjacent in G . We say v is a clockwise neighbor of u at t if v is clockwise to u at t , and counterclockwise neighbor at t otherwise. The width of X is defined to be the quantity

$$w(X) := \min_{v \in V} \max_{u \in V} \delta_X(u, v),$$

which is the length of the shortest path on \mathbb{Z}_n that covers all states of the vertices in the configuration. Let B be a subgraph of G . We denote by $w_B(X)$ the width of the restricted configuration $X|_{V(B)}$ on B .

- (iv) (*locally monotone coupling*) Two coupled oscillators synchronize regardless of the initial configuration, in such a way that the width is a non-increasing function in time.

We say that a discrete system of PCOs with this property is *locally monotone*. It turns out that this local monotonicity turns out to be a very important property to establish one of our fundamental observation. Namely, let τ be a discrete system of PCOs. Note that the dynamic (G, τ_G, X_0) synchronizes if and only if the width $w(X_t)$ converges to 0. Now if τ has locally monotone coupling, then the network on arbitrary connected graph synchronizes once the width becomes small enough ($< n/2$). This also tells us that synchrony in a locally monotone discrete system of PCOs is stable under small perturbation once established.

Lemma 2.2 (width lemma). *Let τ be a locally monotone discrete system of PCOs, $G = (V, E)$ be a graph with a subgraph B , and let X_0 a n -configuration on G . Suppose $w_B(X_0) < n/2$ and no vertex of $G - B$ pulls any vertex of B on time interval $[0, a]$. Then $w_B(X_t)$ is non-increasing on $[0, a]$. Furthermore, the network (G, τ_G, X_0) synchronizes regardless of the structure of G if $w(X_0) < n/2$.*

In fact, similar results for continuous-time models were observed by different authors. For instance a similar observation was used in [18] to show synchronization is guaranteed in an invariant subset. On the other hand, a general theorem in [23] asserts that a similar result applies for systems of continuous-time PCOs where the coupling is of certain type, which includes the type of coupling in our model. Namely, if the initial conditions of the system in [23] are restricted to be $\phi_i(0) \in \{0, 1/n, 2/n, \dots, (n-1)/n\}$, time delay is set to zero and the phase response curve (PRC) is set to $f(x) = -1/n$ for $x \leq 0.5$ and $f(x) = 0$ otherwise, the system exactly recreates the behavior between two n -state oscillators (the difference of whether the blinking state is $n/2$ or 0 is entirely superficial) in our model. However, the model in the paper [23] does not reproduce the exactly same behavior of our discrete model for more than two oscillators, since in their model the coupling effect is *additive* while it is *binary* in ours. Nevertheless, our Lemma 2.2 follows from a nearly identical proof to their theorem. So we only give a brief illustration of the key idea in the case of the firefly networks in Figure 5.

Before we get to Figure 5, we introduce a geometric way of representing discrete systems of PCOs. Usually one visualizes a system of coupled n -periodic oscillators with a circular representation, where each oscillator is represented as a revolving dot on a regular n -gon. However, it turns out that concentrating on *relative states* rather than on the actual states of the oscillators can be useful. More precisely, if the position of an oscillator v is given by a function $X_t(v)$ in time t , then its relative position is given by the function $X_t(v) - t$ modulo n . This can be understood as considering the relative position of each oscillator with respect to an imaginary isolated oscillator, revolving on \mathbb{Z}_n regularly without any interruption. The following is a reformulation of n -periodic firefly networks in terms of relative states.

Definition 2.3 (relative circular representation). *Let $G = (V, E)$ be a graph. Let α be an additional singleton vertex, called the activator. A map $Y : V \cup \{\alpha\} \rightarrow \mathbb{Z}_n$ is called a relative n -configuration on G . Let $S_{\text{rel}}(G, n)$ be the set of all relative n configurations. We say a vertex $v \in V$ is blinking if $Y(v) = Y(\alpha)$ where Y is the current relative configuration. The relative firefly transition map $\tau_G : Y \mapsto Y'$ is defined by*

$$\tau_G(Y)(x) = \begin{cases} Y(x) - 1 & \text{if } x = \alpha \\ Y(x) - 1 & \text{if } x \in V(G) \text{ has a blinking neighbor } u \text{ such that } \delta_Y(x, u) \leq n/2 \\ Y(x) & \text{otherwise} \end{cases}$$

The firefly network $\text{rel}(G, Y_0)$ of period n is the discrete-time dynamical system defined on the relative n -configuration space $S_{\text{rel}}(G, n)$ by the relative firefly transition map τ_G . Suppose an initial configuration Y_0 is given. Then we denote $Y_{t+1} = \tau_G(Y_t)$ for $t \geq 0$.

Note that in the above definition, we have used counterclockwise displacement between two oscillators in relative n -configuration, which is defined similarly as for non-relative configurations (Definition 2.1).

Using terminologies defined above, we can describe the relative firefly transition map as follows: *at each instant, each vertex moves one step counterclockwise if there is a counterclockwise or opposite blinking neighbor, and does not move otherwise.* Geometrically, each blinking node stretches its left arm on the regular n -gon, which is as long as the half of the perimeter, and pulls any of its neighbor within that range. Figure 4 shows a comparison between the standard and relative circular representation of the firefly network. We are going to use those




Figure 4: A comparison between the standard and relative circular representation for $n = 6$, where the relative circular representation uses hexagon. The latter lacks adjacency information between oscillators. For example, in the third configuration, u only pulls v since it is not adjacent to w . If u were adjacent to w , then it would have pulled w as well.

two representations interchangeably.




Figure 5: An example illustrating the proof of width lemma when $n = 8$ and $G = B$. Because of the small width condition on B , the head vertices including x does not affect the tail vertices including w , and all vertices tend to move toward w . The finite size and connection of G will then give the assertion. It is important to note that x never pulls y .

Now we are ready to look at Figure 5, which conveys the key idea behind Lemma 2.2 using the relative circular representation. While the example is for the firefly networks, the same idea was also used in [18] and [23] to prove similar results. Namely, if the initial configuration is concentrated on a sufficiently small arc on the unit circle, then the transition map acts as a non-increasing function on the width; it decreases the width to zero if the underlying graph is connected, in which case we have synchrony. For the firefly networks the initial width should be strictly less than the half, and for other similar models it depends on other parameters such as the length of refractory period or maximum delay(see [18] and [23]).

Next, we discuss a characteristic property of our firefly network. Recall it has locally monotone coupling (iv). As in many other models mentioned earlier, a blinking oscillator may either inhibit its hasty(clockwise) neighbors, or excite its lazy(counterclockwise or opposite) neighbors, or both. In our firefly networks, the coupling is *inhibitory* in the following sense:

- (v) (*inhibitory coupling*) A discrete system of PCOs τ is *inhibitory* if no blinking vertex pulls its counterclockwise neighbor.

This feature of our model enables inductive arguments on certain class of finite graphs by making a small subgraph irrelevant to the dynamic on the rest. For example, consider the example in Figure 5. Since the width is always strictly less than half of the perimeter, for the "head" vertex x , the only neighbor y is too far to inhibit. So x never pulls y , so the dynamic on the triangle with vertices y, z, w is independent of x . Hence if we know something about dynamics on the triangle, we can apply that and we deduce some properties of entire dynamic. This notion of restricting global dynamics on subgraphs is given below precisely.

Definition 2.4. Let $G = (V, E)$ be a graph and $X_0 : V \rightarrow C_n$ be a n -configuration on G . Let H be a subgraph of G . Let τ be the firefly transition map as given in Definition 1. We say the dynamic (G, X_0) restricts on H if the restriction and transition maps commute, i.e., $\tau_G(X_t)|_H = \tau_H(X_t|_H)$ for all $t \geq 0$. We say the dynamic (G, X_0) restricts on H eventually if there exists $r \geq 0$ such that (G, X_r) restricts on H .

3. The branch width lemma, 1-branch pruning, and the path theorem

In this section we establish a key lemma on n -periodic firefly networks, which will play a central role in the proof of Theorem 2 and Theorem 4. Our argument for the path theorem goes as follows. Let P be a path of length m and let $n \geq 3$ be period. First note that the lower bound is the time that it takes to synchronize a configuration where all vertices have state $n - 1$ but one end vertex, which has the blinking state $b(n)$ (See Figure 6).

For the upper bound, starting with any initial configuration on a finite path, we will observe that the two end vertices of the path do not affect their neighbors eventually, so they become irrelevant to the dynamic of the rest. Hence we can apply inductive arguments on a smaller sub-path. To get this restriction property on paths, we analyze induced local dynamics on the last two vertices of a path, which we call a 1-branch. The general notion of k -branch is given below, which will serve a similar role for the tree theorems.

Definition 3.1. Let $G = (V, E)$ be a graph. A connected subgraph $S \subseteq G$ is called a k -star if it has a vertex v , called the center, such that all the other vertices of S are leaves in G . A k -star S is called a k -branch, if the center of S has only one neighbor in $G - S$. We may denote a k -branch by B rather than by S .




Figure 6: An example of 6-periodic network on a path of length $m = 8$ which takes $6 \cdot (6/2 - 1 + 8)$ seconds to synchronize. The "color difference" of $n/2$ on the first edge is distributed after $n(n/2 - 1)$ seconds, and then each "unit difference" moves 1 step right after every n seconds. So we need nm more seconds to take away all particles and achieve synchrony.

Notice that in the definition above, not only S is isomorphic to the canonical star graph, but also the leaves of S must be leaves in G . Now the following lemma establishes this restriction property of end vertices of finite paths.

Lemma 3.2. [1-branch pruning lemma] *Let $G = (V, E)$ be a graph with a vertex w . Suppose there is a 1-branch B rooted at w with center u and leaf v . Let $n \geq 3$ and let X_0 be a n -configuration on G . Let $H = G - v$. Then we have the followings:*

- (i) *The dynamic (G, X_0) restricts on H eventually;*
- (ii) *Suppose H is n -synchronizing. Then G is n -synchronizing. Furthermore, if H synchronizes every n -configurations in $N_H \geq 0$ seconds, then G synchronizes every n -configurations in $n^2/2 + 2n - 2 + N_H$ seconds.*




Figure 7: A graph G with a 1-branch B

Notice that Theorem 2 follows immediately by an induction, where locally monotone coupling (condition (iv) in Section 2) gives the base case and the above lemma gives the induction step.

Before we look into details, we discuss an interesting applications of the theorem. Imagine we have achieved synchrony of oscillators on a graph G , and we wish to add a new oscillator w to the already-synchronized network. Notice that this is to consider the dynamic on $G + w$ where all vertices of G have the same state and w may have an arbitrary state. The new dynamic is guaranteed to synchronize by the width lemma (Lemma 2.2) when n is odd, since then any configuration on $G + w$ with only two states would have width $< n/2$. However, it is possible that the width is exactly $n/2$ when n is even. The following corollary of Theorem 2 implies that we are guaranteed to obtain synchrony in this case too.

Corollary 3.3. *The firefly network $\mathbf{re1}(G, Y_0)$ synchronizes regardless of G if the relative initial configuration consists of only two states.*

This corollary follows from the following simple but interesting observation; an initial configuration with only two states gives a "path decomposition" of G , i.e., a quotient map of G onto a path, and the dynamics on G is essentially that on this path; so it is guaranteed to synchronize by Theorem 2. This is illustrated in Figure 8.




Figure 8: An example of 6-periodic network illustrating Corollary 3.3. Recall that $b(6) = 2$ is the blinking state. In the above example there are only two states 2 and 5 initially. Let A and B be the set of vertices with initial state 2 and 5, respectively. Let $A_0 \subseteq A$ be the neighbors of B , and let $A_1 \subseteq A \setminus A_0$ be the neighbors of A_0 , and let $A_2 \subseteq A \setminus A_0 \setminus A_1$ be the neighbors of A_1 , and so on. We partition B similarly. In the example above, the vertices at the same "level" represents each class A_i and B_j . Then we identify each class into single vertex. This quotient gives a path, and the network on G behaves like that on this path.

The following lemma gives the key observation in this paper. It says that if a graph has a branch, and if the initial configuration has a small width on the branch, then the leaves of the branch become irrelevant to the dynamics on the rest eventually. This is called the *branch width lemma*, which is a variant of the width lemma (Lemma 2.2) using a specific inhibitory structure of the coupling in our model.

Lemma 3.4 (branch width lemma). *Let $G = (V, E)$ be a graph with a vertex w . Suppose there is a k -branch B rooted at w , with center v and leaves l_1, \dots, l_k , $k \geq 1$. Let H be the graph obtained from G by deleting the leaves of this branch. Let $n \geq 3$ and let X_0 be a n -configuration on G . Suppose $w_B(X_0) < n/2 - 1$. Then we have the followings:*

- (i) v is clockwise to all leaves of B at some time $r \leq n(w_B(X_0) + 1)$ and $w_B(X_r) \leq w_B(X_0)$;
- (ii) v is clockwise to all leaves of B for all $t \geq r$, and $w_B(X_t) \leq w_B(X_0) + 1$ for all $t \geq 0$;
- (iii) If v is clockwise to all leaves at $t = r$, then the dynamic (G, X_r) restricts on H ;
- (iv) If H is n -synchronizing, then (G, X_0) synchronizes.




Figure 9: A graph G with a k -branch B

A detailed proof of this lemma is given in the Appendix A, and here we give a quick illustrative explanation. Suppose $n = 8$ and $k = 3$. Since the coupling is inhibitory, the three leaves of B and the root w only inhibit the center v , and during this period, Lemma 2.2 (i)

keeps the small width on the leaves of B . So eventually, we will have a situation as in the first diagram in Figure 10, where the branch width $w_B := w|_B$ is still strictly less than $n/2 - 1$ and the center v is at the "tail". Now the root w will pull v occasionally, increasing the branch width by 1. But since we have a wiggle room on the branch width, the increased branch width is still small ($< n/2$) and the leaves do not pull v until they blink again. Then the center blinks and pulls all the leaves, decreasing the branch width by 1. Hence the original branch width is recovered, and because of the small width on the branch, the leaves never pull the center.




Figure 10: An illustration of branch width recovery for $n = 8$ and $k = 3$. Once v is at the tail, w can pull v to increase the branch width by 1 but v pulls the head leaf and decrease the branch width by 1, before w blinks again. Note that w can get external pullings from its neighbors different from v but it doesn't affect our argument.

Note that the 1-branch pruning (Lemma 3.2) follows easily from the above lemma. Indeed, let G, w, v , and u be as in Lemma 3.2. It suffices to show that eventually, the width on this 1-branch, which is just the minimum of $\delta_t(u, v)$ and $\delta_t(v, u)$, becomes less than the threshold $n/2 - 1$. This is illustrated in Figure 11, and the upper bound in Lemma 3.2 (ii) can be computed easily from the figure.




Figure 11: (a) If $\delta_0(v, u) \leq n/2$ when u first blinks, then $\delta_t(v, u)$ monotonically decreases to zero. (b) Otherwise $\delta_0(u, v) < n/2$ when v first blinks, and the root w does not pull the center v until the activator moves from u to v , and v pulls u to make $\delta_t(u, v) < n/2 - 1$.

4. Local configurations, Poincaré return map, and tree theorems

In this section we discuss firefly networks on trees. More specifically, we prove Theorem 3 and Theorem 4 for $n \in \{3, 4, 5\}$. The bottom line of our proofs is the following. Let T be a tree and X_0 be a n -configuration for some $n \in \{3, 4, 5\}$, and suppose for contrary that (T, X_0) is a minimal counterexample to Theorem 4. Then an analysis of the induced local dynamic on a particular branch of T would yield that the global dynamic must restrict onto a smaller subtree, which contradicts to the minimality.

An important concept for such local analysis is Poincaré return map, which is to take "snapshots" of system configurations where a fixed special agent of the system has a particular state, and look into the induced dynamic on the set of such snapshots. This technique effectively reduces dimensionality of system configurations, and was used by Mirollo and Strogatz [21] by taking snapshots only when a particular oscillator blinks. In this paper, we

incorporate similar technique to analyze induced local dynamic on a branch, where the center of a fixed branch is taken to be the special vertex.

Definition 4.1. Let $G = (V, E)$ be graph, $n \geq 3$, and let Y be a relative n -configuration on G . Let v be a vertex, $N \subset N_G(v)$ be a subset of all neighbors of v , and let $N_v := N \cup \{v\}$. Then the local configuration of X on N_v is the restriction $Y|_{N_v}$. If $N_v = N_G(v)$, then we write $Y^v := Y|_{N_v}$. Fix a relative initial n -configuration Y_0 . We say that with respect to the dynamic $\text{rel}(T, Y_0)$, a local configuration on N_v is recurrent if it occurs infinitely often, and transient otherwise.

The following observation quickly gives us some transient local configurations at some vertex v with a leaf.

Lemma 4.2 (opposite leaf lemma). Let $G = (V, E)$ be a graph with a leaf u and its neighbor v . Let Y_0 be any relative initial configuration on G . Suppose Y^v is any relative local configuration at v where v blinks and the counterclockwise displacement $\delta_t(u, v)$ is $\lfloor n/2 \rfloor$. Then Y^v is transient.

Proof. This proposition can be best understood by looking at the following cases when $n = 6$ in Figure 12. Suppose the first local configuration in Figure 12 (a) is recurrent at the 1-branch $B := v + u$ for some leaf u of v . We may back-track for 3 seconds. During this period the center v does not blink so the leaf u is not pulled, and v couldn't have been pulled by any of its neighbors for the first 2 backward iterations, since it was counterclockwise to the activator during that period. Hence the local configuration before three seconds must have been the fourth configuration in Figure 12 (a), but the first iteration from this conflicts to the pulling of u on v . Thus the shaded local configuration in Figure 12 (a) is transient. Since this holds for all 1-branch centered at v , the assertion for the n = even case follows. Figure 12 (b) for $n = 5$ illustrates the similar argument for the odd period case. \square




Figure 12: An illustration of the opposite leaf lemma for (a) $n = 6$ and (b) $n = 5$. The square and dot represents the relative states of the center v and leaf u , respectively. The two shaded local configurations are transient.

We call a local configuration $Y|_B$ on a branch B *opposite* if $Y|_{B'}$ is such a local configuration in the above proposition for some $B' = v + u$, where v is the center of B and u is some leaf of B ; we call $Y|_B$ *non-opposite* otherwise.

We start proving Theorem 4 for $n = 3$. Here we only need Lemma 3.2 and Lemma 4.2 for local analysis.

Proof of Theorem 4 for $n = 3$. The "only if" part is trivial. Let us show the "if" part. Let $T = (V, E)$ be a tree, and let $X_0 : V \rightarrow \mathbb{Z}_3$ be a 3-configuration such that every vertex of T blinks in the dynamic (T, X_0) . We wish to show that (T, X_0) synchronizes. We use an induction on $|V|$. Since K_2 is 3-synchronizing, we may assume $|V| \geq 3$. If T has a 1-branch, then we are done by the induction hypothesis and Lemma 3.2. So we may assume there is no 1-branch. Since

we can identify any two leaves of the same state with common neighbor, we may assume all leaves of each k -star in T have distinct states all time.

Now let u be a leaf in T , and let v be the neighbor of u . Since $|V| \geq 3$ and $T - u$ is connected, v has a neighbor w in $T - u$. Let B be a k -branch consisting of v and all of its leaf neighbors. Note that $k \geq 2$ by our assumption. By the hypothesis the center v blinks infinitely often, so we can apply Lemma 4.2, which says in this case that eventually, whenever v blinks (has state 1), all leaves of v must have either state 1 or 2. Thus we may assume that v has exactly two leaves which eventually have state 1 and 2 whenever v blinks. But at any such instant, v pulls the leaf of state 2 and synchronizes its two leaves. This contradicts our assumption that no two leaves of v ever have the same state. This shows the assertion. \square

To proceed more concisely in proving Theorem 4 for $n = 4$ or 5 , we consider the following class of configurations for which our inductive argument may not work.

Definition 4.3. Let $G = (V, E)$ be a graph. A n -configuration $X_0 : V \rightarrow \mathbb{Z}_n$ on G is irreducible if the dynamic (G, X) never restricts to a proper subgraph of G of at least 2 vertices. If X_0 is irreducible, we say (G, X_0) is an irreducible dynamic.

Notice that if (G, X_0) synchronizes, then X_0 cannot be irreducible, since after the synchrony the dynamic can be restricted on any proper subgraph of G . The next proposition tells us that any initial configuration for a minimal counterexample for Theorem 4 is irreducible if every vertex blinks infinitely often in the dynamic.

Proposition 4.4. Let $T = (V, E)$ be a tree, and suppose that there exists a n -configuration $X_0 : V \rightarrow \mathbb{Z}_n$ such that every vertex blinks infinitely often in (T, X_0) but the network does not synchronize. Further assume that T is a smallest such tree. Then X_0 is irreducible.

Proof. Suppose not. Then there exists a proper subtree T' of T with ≥ 2 vertices and an integer $r \geq 0$ such that the network (T, X_r) restricts on T' . Then every vertex of T' blinks infinitely often in $(T', X_r|_{T'})$ since they do so in the larger network (T, X_0) . Then by the minimality of T the restriction synchronizes. After the synchrony on T' , we may identify the vertices of T' . In other words, eventually, we can contract T' to a single vertex without affecting the dynamic. Denote the resulting graph by T/T' , which is a proper minor of T . So T/T' is a tree which is strictly smaller than T . Clearly every vertex of T/T' blinks infinitely often in the induced dynamic, so we get a synchrony on T/T' . But this tells us (T, X_0) eventually reaches synchrony, which is a contradiction. Therefore X_0 must be irreducible. \square

Hence to obtain Theorem 4 for period 4 and 5, it suffices to show that if there is a minimal counterexample (T, X_0) then X_0 cannot be irreducible. In the proof of $n = 3$ case, we used the fact that a certain local configuration on a branch is transient, and then Lemma 3.2 to restrict the global dynamic on a smaller subtree.

Proposition 4.5. Let $T = (V, E)$ be a tree, $n \geq 3$, and let X_0 be a n -configuration which is not irreducible. Suppose T has a k -star S and let v be its center. Then any local configuration $Y|_S$ is transient if it satisfies either of the following conditions:

- (i) Some two distinct leaves of B have the same state in Y ;
- (ii) $Y|_S$ is opposite.

Furthermore, if $S = B$ is a branch, then $Y|_B$ is also transient if it satisfies either of the following conditions:

- (iii) Only a single state is used for the leaves of B ;
- (iv) $Y|_B$ has width $< n/2 - 1$.

Proof. For (i), suppose that two leaves x, y of B have the same state at some point. Then they will always have the same state later on, so we may restrict the global dynamic onto $T - x$, for example. This contradicts the irreducibility. Lemma 4.2 shows (ii) is transient. Now suppose $S = B$ is a branch. Then the 1-branch pruning lemma (Lemma 3.2) shows (iii) is transient. Finally, (iv) follows from the branch width lemma (Lemma 3.4). \square

Proposition 4.6. Let $G = (V, E)$ be a graph with an induced k -star S centered at a vertex v . Let $n \in \{4, 5\}$, and suppose X_0 is an irreducible n -configuration on G such that v blinks infinitely often in the dynamic. Then we have the followings:

- (i) If $n = 4$, then all local configurations on S where w blinks are transient except Figure 13 a;
- (ii) If $n = 5$ and if $S = B$ is a branch with root w , then all local configurations on B where v blinks are transient except Figure 13 b.




Figure 13: Possibly recurrent local configurations on stars for $n = 4$ and on branches for $n = 5$.

Proof. For $n = 4$ it follows easily from Proposition 4.8 (i) and (ii). We give a detailed argument for $n = 5$ case. Suppose S is a branch B rooted at w . So the neighbors of v are the leaves of B and the root w . By the irreducibility, we may assume that no two leaves ever have the same state. Consider the local configurations on S where v blinks. First of all, there are 15 non-opposite local configurations on S as in Figure 14.




Figure 14: Non-opposite local configurations on a branch B with some of the possible transitions between them. The dots represent leaves, and the square represent the center of the branch. If the configuration is irreducible, only the shaded one is recurrent.

Notice that the 6 local configurations on the right are transient by (iii) and (iv) of the previous proposition. Furthermore, the first five local configurations in the second row lead to some of the 6 local configurations on the right, after the center v pulls some of its leaves. Hence they are also transient.

Now to rule out the first three in the first row, we consider all possible transitions from them. Remember that v could be pulled by the root w during this transition, so the next position of v is uncertain in the local sense. But note that after the center v pulls the leaves ahead

of itself, there is no further change on the leaves until v blinks again, so the next position of leaves are determined locally. Hence the first three in the first row can only lead to the second and third in the second row, which are transient. But since v blinks infinitely often, at least one local configuration where v blinks must be recurrent. Thus the assertion follows. \square

Now Theorem 4 for $n = 5$ follows directly from the following lemma:

Lemma 4.7 (branch pruning for $n = 5$). *Let $G = (V, E)$ be a graph with a vertex w . Suppose there is a k -branch B rooted at w with center v (see Figure 9). Then there is no irreducible 5-configuration on G such that every vertex blinks infinitely often in the dynamic.*

Sketch of proof. Let X_0 be any 5-configuration on G where all vertices blink infinitely often in the dynamic. It suffices to show that the dynamic on G restricts on $G - B$ eventually. For this we show that the center v does not pull w eventually. By Proposition 4.6, B has two leaves which have states 1 and 4 whenever v blinks eventually. This specific dynamic on B is possible only if w blinks at some specific moments, and an easy analysis shows that this yields w must have state 0 whenever v blinks (a similar analysis is given for the $n = 4$ case). Hence w is not pulled by v eventually, as desired. \square

It now remains to show Theorem 4 for $n = 4$. Interestingly, there is no 4-periodic counterpart of Lemma 4.7. That is, not every branch can be pruned out under the assumption of irreducible configuration where all vertices blink infinitely often. However, this pruning is possible when the graph is a tree and the branch is at the end of a longest path.

Lemma 4.8. *Let $G = (V, E)$ be a graph. Suppose that G has an irreducible 4-configuration X_0 such that every vertex of G blinks infinitely often in (G, X_0) . Suppose G has a k -star S for $k \geq 2$ with center v . Then the followings are true:*

- (i) $k = 2$ and whenever v blinks (has state 1), its two leaves have state 0 and 1.
- (ii) If $S = B$ is a branch in G , then in the limit cycle the local dynamic on B and its root repeats the following sequence:

$$\begin{array}{c|cccccccccc} \text{center} & \text{leaves} & 1|01 & 2|12 & 2|23 & 3|30 & 3|01 & 3|12 & 3|23 & 0|30 & 1|01 \\ \text{root} & & 3 & 3 & 0 & 1 & 2 & * & * & * & 3 \end{array} \quad (3)$$

(iii) In case of (ii), the root of B must have at least three neighbors in $G - B$.

(iv) There are no two branches rooted at the same vertex.

Proof. (i) directly follows from Proposition 4.9. Now suppose S is a branch with root, say, w . For the other parts, we analyze the actual transition between the unique recurrent local configuration a in Figure 13. The analysis in Figure 15 shows that starting from this local configuration, the center v must be pulled by the root w at the third or fourth iteration, and each transition takes exactly 8 seconds.

Note that In the standard representation, the recurrent local configuration a in Figure 13 can be represented as $1|01$, by which we mean state of center|state of leaves. Hence the above analysis shows that $1|01$ must lead to $3|01$ after 4 seconds. We consider the first 4 iterations starting from $1|01$ in standard representation:

$$\begin{array}{c|ccccc} \text{center} & \text{leaves} & 1|01 & 2|12 & *|23 & *|30 & 3|01 \\ \text{root} & & x_1 & x_2 & x_3 & x_4 & x_5 \end{array} \quad (4)$$




Figure 15: Actual transition between the unique recurrent local configuration a in Figure 13 when $S = B$ is a branch. The center is pulled by one of its leaves in 2 seconds, and for the next 2 seconds it may or may not be pulled by the root w , so it splits into two cases. The upper one is for no pulling from w , which ends up with a wrong local configuration. Hence the center v must be pulled by w at third or fourth iterations. On the other hand, the dashed arrow to the third column indicates a transition where v is pulled by w . Since w blinks at most once in every 4 seconds, v can be pulled by w at most once in every 4 seconds. Hence there is no external pulling in the following 2 seconds. Then the case splits into two as before, and this time w should not pull the center to end up with the correct local configuration. The whole transition takes exactly 8 seconds.

where $1 \in \{x_3, x_4\}$. However, $x_3 = 1$ yields $x_2 = 0$ and $x_1 = 3$ which is impossible since the blinking center then must have been pulled $x_1 = 3$ so that $x_2 = 3$. Therefore $x_4 = 1$, and consequently, $x_1 = 3$. Thus the 8 iterations and each configurations on the branch and root is given by (9). This shows (ii).

For (iii), observe that in the sequence (9), the root w must be pulled 3 times in the last four seconds. Since v do not pull w during this iterations, and since w can get at most one pulling from each of its neighbor in every 4 seconds, we conclude that w must have at least 3 neighbors different from v . This shows (iii).

To show (iv), suppose there are two branches B_1 and B_2 rooted at w . By (ii), the dynamics on those branches in the limit cycle are given by (9). Note that none of $*$'s in (9) is 1 for trivial reason. Now whenever w blinks, the centers and leaves in both branches must have the same configuration, namely, 3|30. This determines the dynamics on those branches completely, and we see that they must have the identical dynamic. Thus the dynamic on G reduces to $G - B_1$, which contradicts the irreducibility. This shows (iv). \square

Proof of Theorem 4 for $n = 4$. It suffices to show the "if" part. Let (T, X_0) be a minimal counterexample to Theorem 4 for $n = 4$. It suffices to show that X_0 is not irreducible. Suppose the contrary. We may assume $|V| \geq 3$. We first show that T must have an "inevitable" local structure as in Figure 16. Let P be a *longest* path in T . It has length at least 2 since $|V| \geq 3$. Let u be an end of P , v the neighbor of u in P , and w the neighbor of v in $P - u$. Then by the choice of v and Lemma 4.8 (i), v and its leaves form a 2-branch rooted at w . Let us denote this branch by B . By Lemma 4.8, the dynamic on B and the root w repeats the sequence (9). By Lemma 4.8 (iv), there is no other branch rooted at w . It then follows from our choice of w from a maximal path P , that every neighbor of w not in P is a leaf. We conclude that w must have at least two leaves by Lemma 4.8 (iii). Then by Lemma 4.8 (i), w and its leaves form a 2-star.




Figure 16: A local structure of T assuming (T, X_0) is a minimal counterexample for Theorem 4 for $n = 4$

Now by Lemma 4.8 (i), eventually, the two leaves of w must have state 0 and 1 whenever v blinks. Now we insert the states of leaves of w into the sequence (9). The first four iterations would then be as follows:

$$\begin{array}{l|l} u & \text{leaves of } u \\ w & \text{leaves of } w \end{array} \quad \begin{array}{cccccc} 1|01 & 2|12 & 2|23 & 3|30 & 3|01 \\ 3|12 & 3|23 & 0|30 & 1|01 & 2|12 \end{array} \quad (5)$$

Recall that the above sequence describes four iterations starting from an arbitrary instant when v blinks. Notice that there is a blinking leaf of w at the first instant. Thus we conclude that whenever v blinks w has a blinking neighbor so that the pulling of v on w is redundant. Thus the dynamic on T restricts on $T - B$ eventually, contrary to our assumption. This shows the assertion. \square

We have seen in the introduction that Theorem 3 follows directly from Theorem 4 and Lemma 5. The proof of Lemma 5 is independent from our storyline, so we give it in Appendix B.

5. Randomized self-stabilization of firefly network on arbitrary connected graphs

Let τ be a n -periodic discrete system of PCOs for some $n \geq 3$. In the case of firefly networks, we have seen that non-synchronizing limit cycles depend heavily on the symmetry of the graph and configurations on it. In order to break this symmetry to enhance synchronization, we introduce randomness to our GCA model in the following way. For a given connected graph $G = (V, E)$, consider an edge weighting $\epsilon : E \rightarrow (0, 1)$. Now each edge e in G will be present with a fixed probability $\epsilon(e)$ independently at each instance in the dynamic (G, τ_G) , so two adjacent oscillators along the edge e will now see each other with the associated probability. In other words, we are introducing stochastic reception of pulses. This introduces randomness to our deterministic GCA model and turns it into a Markov chain, which we denote by a triple (G, ϵ, τ_G) . We identify all constant configurations into a single state called *sync*, and take

$$\mathcal{S} := \{X : V \rightarrow \mathbb{Z}_n \mid X \text{ is non-constant}\} \cup \{\text{sync}\}$$

as the state space of our Markov chain. Now clearly *sync* is an absorbing state, meaning that once the system state is *sync*, then it is so thereafter. In fact, *sync* is the only absorbing state in this Markov chain. Given that, Theorem 6 follows from an elementary argument in Markov chain theory.

Showing that *sync* is the unique absorbing state is based on the following simple observation. Note that it suffices to show that for any given initial configuration X , there is a positive probability $p_X > 0$ such that the Markov chain enters *sync* eventually. To this end, suppose we have achieved synchrony on a connected subgraph $H \subset G$. By the connection of G , we can pick a vertex w in $G - H$ that is adjacent to some vertex of H . Since the states on H has been synchronized, there are only two states on $H + w$. So if all edges in $H + w$ are present and the bridging edges between $H + w$ and the rest are absent for a sufficient amount of time, $H + w$ behaves as a deterministic network isolated from the rest, so we have synchrony on it by Corollary 3.3. We repeat this process until we expand a partial synchrony to the entire G . Hence one can synchronize arbitrary configuration X with some positive probability through this process. Note that the time that it takes for each step is bounded above by the upper

bound given in Theorem 2. Then an elementary calculation gives an upper bound on the expected time until absorption. However, the upper bound that we get from this argument does appear to be far from optimal.

The above argument works for any discrete system of PCOs for which the two-state corollary (Corollary 3.3) holds. While a general locally monotone discrete system of PCOs may lack this property, it is easy to modify this argument to work for arbitrary locally monotone discrete system of PCOs, by using Lemma 2.2 instead of Corollary 3.3, which holds for any locally monotone discrete system of PCOs. Moreover, we may give the stochasticity not by assigning a fixed probability to be absent for each edge, but for each vertex, as in [18]. A similar argument easily applies, so we obtain the following generalization of Theorem 6:

Theorem 5.1. *Let τ be a n -periodic locally monotone discrete system of PCOs for any $n \geq 3$ and let $G = (V, E)$ be a connected graph. We give stochasticity by assuming either of the followings:*

- (i) (stochastic emission) *Each vertex $v \in V$ is present independently at each instant with a fixed probability $p_v \in (0, 1)$;*
- (ii) (stochastic reception) *Each edge $e \in E$ is present independently at each instant with a fixed probability $p_e \in (0, 1)$.*

Then the Markov chain (G, τ_G) is an absorbing chain with sync being the unique absorbing state. In particular, every n -configuration on G synchronizes with probability 1 in finite expected time.

6. Concluding remarks

We defined a GCA model for pulse-coupled oscillators and studied their network behavior, mainly focused on various conditions for synchrony. Taking the advantage that the dynamic of individual oscillators is extremely simple, we were able to obtain conditions on initial configurations and network topologies that guarantee synchronization. Paths were generic to our model in the sense that every finite path is n -synchronizing for all $n \geq 3$. We then studied to what extent this self-stabilization property on paths extends, and we obtained a local-global principle on tree networks for period $n \leq 6$ (a proof for $n = 6$ case was omitted). We also showed that any n -periodic firefly network on random networks synchronizes with high probability starting from an arbitrary initial configuration, where random networks are given by introducing independent random errors either to the vertices or edges of connected graphs.

The remainder of this section contains a description of directions of future research. Our first cornerstone was Theorem 2. An obvious extension of this result is to consider infinite path, namely, the integer lattice \mathbb{Z} , instead of finite paths. From the example in Figure 6, however, it is apparent that one can have a non-synchronizing initial configuration on \mathbb{Z} for any period n ; e.g., a "wave of pullings" could propagate from $-\infty$ to ∞ . An appropriate type of question to ask in this case might be as follows: starting from a uniform product measure on \mathbb{Z} , and fixing a finite interval I , does one have synchronization on I with probability 1? If so, how would such a probability scale in time?

For low periods $n \leq 6$, we have seen that the firefly transition map synchronizes every initial configuration on an arbitrary fixed tree, given that the maximum degree is less than the period. However, we have seen that Theorem 4 is not valid for $n = 7$ in Figure 3, and verifying the theorem for $n \geq 8$ is open. On the other hand, we have also seen that cycles

and cliques are not synchronizing in general. Indeed, Dolev [10] discusses that given any distributed synchronization algorithm, one can always find a non-synchronizing configuration on some cycle. Hence, one might think that not containing cycle as a subgraph is a critical factor for a network to be synchronizing. However, K_3 is 6-synchronizing and the "shovel graph", obtained by vertex-summing K_3 at the end of a path, is also 6-synchronizing. Thus just having a cycle in the network does not necessarily mean that the network is not 6-synchronizing. It is a future goal to obtain a complete characterization of n -synchronizing graphs. Expanding the class of n -synchronizing graphs is of special interest in the view point of self-stabilizing networks, since it would allow us to design a fault-tolerant and self-synchronizing system with variety of network topologies.

In section 5 we incorporated a certain stochasticity into our deterministic model and obtained absorbing Markov chains with synchrony being the unique absorbing state. Having established this universal synchrony with high probability, it is then interesting to ask what is the expected time until synchrony. While for each graph G we can get exact expected time until absorption using an elementary Markov chain theory, the recursive argument using Corollary 3.3 or Lemma 2.2 gives a trivial upper bound for the expected time until synchrony, which only depends on the number of vertices, maximum degree of G , and the period. However, this trivial upper bound is too crude to be precisely stated in this paper. On the other hand, take $G = K_m$ and assume stochastic reception with $p_e \equiv p(n) \in (0, 1)$. Then the resulting Markov chain is a firefly network on a sequence of Erdős-Renyi random graphs $G(m, p)$. It would be interesting to study phase transitions in this model with respect to the order of p .

Finally, one can also study different transition maps, in comparison with our firefly transition map. For any class of connected graph \mathcal{H} , call a transition map τ on the space of n -configurations \mathcal{H} -type if it synchronizes every n -configuration on all $H \in \mathcal{H}$. Our firefly transition map, for example, is a path-type which synchronizes tree networks for low periods given a degree condition. One can then study \mathcal{H} -type transition maps for various choices of \mathcal{H} . Recall that, somewhat contrary to our intuition, the all-to-all networks were not synchronizing in general; K_4 is not 6-synchronizing as we have seen in Figure 1 (d) in the introduction, and the $2n + 1$ -configuration using states 0, n , and $2n$ on K_3 does not synchronize for all $n \geq 2$. So any clique-type transition map, if any, must be essentially different from our firefly transition map. However, notice that by Dolev's argument \mathcal{H} cannot contain the class of cycles. Studying different types of transition maps should be useful for various application of different nature.

A. Proof of Lemma 3.4

Proof of Lemma 3.4. Observe that if any two leaves have the same state at some point, then they will always have the same state since they get same input from u and they have no other neighbors. Hence we may identify any two leaves after they have the same state. Suppose X_0 is an n -configuration on G with $w_B(X_0) < n/2 - 1$. Call a vertex of B at the head(tail) at time t if it is counterclockwise(clockwise) to all the other vertices of B at time t .

We first show (i). If u is at the tail, clearly we may take $r = 0$. Otherwise, there exists some leaf at the tail, say l_k without loss of generality, which is strictly clockwise to u . Observe that both the root w and any leaf strictly clockwise to u , if any, will pull u towards the tail. So there will be some $r \geq 0$ such that u is at the tail at time r . We may take r as small as possible. On the time interval $[0, r)$, Lemma 2.2 tells us that $w_B(X_t)$ is non-increasing as long as w does not

affect u , but when w does affect u , $w_B(X_t)$ would not increase since l_k is strictly clockwise to u . Thus we have $w_B(X_r) \leq w_B(X_0)$. Furthermore, we have the upper bound $r \leq n(w_B(X_0) + 1)$ since it takes $\leq n$ seconds for l_k to blink for the first time and it needs to pull u by $\leq w_B(X_0)$ steps, and l_k blinks once in every n seconds on time interval $[0, r]$. This shows (i). Without loss of generality, let l_1 be leaf at the head at time $t = r$ (see Figure 10).

Next we show (ii). By part (i), It suffices to show that for all $t \geq r$, u is clockwise to all leaves of B and $w_B(X_t) \leq w_B(X_r) + 1$. The idea is that if w pulls u to increase the branch width by 1 after time r , then u will pull the head leaf l_1 before it gets pulled by w once more and recovers the original branch width (see Figure 7). If w never pulls u for $t \geq r$, then the assertion follows immediately. So suppose at some time $t \geq r$ that w blinks and u is counterclockwise to w . We are going to see what will happen until w blinks again. u will be pulled toward w , increasing the branch width by 1, but which is $\leq w_B(X_0) + 1 \leq n/2$. So the branch width is non-increasing until w blinks again. In fact, u will blink before w does so, and it will pull all the strictly clockwise leaves, in particular, l_1 , and decrease the branch width by 1. Note the branch has recovered its original branch width and u is still at the tail. The assertion is clear if w never blinks again, and otherwise we are back to the previous case. This shows (ii).

By the hypothesis and (ii), we have $w_B(X_t) < n/2$ for all $t \geq 0$ and u is at the tail for all $t \geq r$. Hence no leaves of B pulls u on time interval $[r, \infty)$, so the dynamic (G, X_r) restricts on H . If H is n -synchronizing, this means eventually, all vertices of H will have the same state as u . Then the branch width becomes the total width, and we have synchrony on G by Lemma 2.2. This shows (iii) and (iv). \square

B. Proof of Lemma 5

Proof of Lemma 5. Suppose for contrary that u does not blink after some $t_0 \in \mathbb{N}$. Consider the clockwise displacement $\delta_t(u, \alpha)$ from u to the activator α . Note that by definition, u blinks at t if and only if $\delta_t(u, \alpha) = 0$. Since u does not blink after $t = t_0$, the displacement is positive for $t > t_0$. Note that the displacement is non-increasing until the activator catches up to u , and it strictly decreases whenever no neighbor of u is blinking, since then u does not move while α moves toward u . Hence we may assume at $t = t_1 \geq t_0$, the displacement stabilizes to its minimum and at least one neighbor of u blinks at each second $t \geq t_1$.

Since the degree of u is at most $n - 1$, there is always an unoccupied state on \mathbb{Z}_n by the neighbors of u . Suppose there is an empty spot with counterclockwise displacement $\leq n/2$ from the activator as in Figure B.1 (a). Then no neighbor of u whose state is between the empty spot and the activator moves while the activator proceeds toward the empty spot, so when the activator occupies the empty spot, there is no blinking neighbor of u , contrary to our assumption. Hence each state with counterclockwise displacement $\leq n/2$ from the activator must be occupied by at least one neighbor of u for all $t \geq t_1$. Now let $h(t)$ be the number of neighbors of u that are strictly behind the activator at time $t + t_1$, that is, that are with clockwise displacement $< n/2$ from the activator. Note that any oscillators that were blinking at $t = t_1$ fall behind the activator at $t = t_1 + 1$, so $h(1) \geq 1$ (see Figure B.1 (b)). Note also that this neighbor strictly behind the activator will be counted by the function h in the next $b(n) - 1$ seconds. As the activator proceeds one more step counterclockwise, at least one more oscillator falls behind the activator and we have $h(2) \geq 2$. Suppose $n = 2m$. The similar argument gives $h(m - 1) \geq m - 1$. But by the assumption all the $m + 1$ spots within counterclockwise displacement $\leq n/2$ from the activator must be occupied by the neighbors of u . It follows that at $t = t_1 + m - 1$,

there are at least $2m = n$ neighbors of u , a contradiction. The similar argument leads to a contradiction for the case $n = 2m + 1$. Thus u must blink infinitely often.




Figure B.1: (a) The activator occupies any forward empty spot without pulling any neighbor of u to the spot. (b) More oscillators fall behind the activator but the states counterclockwise to the activator must be all occupied. * indicates the spots that may or may not be occupied by neighbors of u .

□

Acknowledgement

We give special thanks to David Sivakoff for his priceless advices, and also to Woong Kook and Hyuk Kim for their encouragement on this work. In addition, we appreciate the referees for their insightful comments and suggestions.

References

- [1] Acebrón, J. A., Bonilla, L. L., Vicente, C. J. P., Ritort, F., Spigler, R., 2005. The kuramoto model: A simple paradigm for synchronization phenomena. *Reviews of modern physics* 77 (1), 137.
- [2] Antsaklis, P., Baillieul, J., 2004. Guest editorial special issue on networked control systems. *Automatic Control, IEEE Transactions on* 49 (9), 1421–1423.
- [3] Antsaklis, P., Baillieul, J., 2007. Special issue on technology of networked control systems. *Proceedings of the IEEE* 95 (1), 5–8.
- [4] Arora, A., Dolev, S., Gouda, M., 1992. Maintaining digital clocks in step. In: *Distributed Algorithms*. Springer, pp. 71–79.
- [5] Buck, J. B., 1938. Synchronous rhythmic flashing of fireflies. *The Quarterly Review of Biology* 13 (3), 301–314.
- [6] Bullo, F., Cortés, J., Piccoli, B., 2009. Special issue on control and optimization in cooperative networks. *SIAM Journal on Control and Optimization* 48 (1), vii–vii.
- [7] Chopard, B., Droz, M., 1998. *Cellular automata modeling of physical systems*. Vol. 24. Cambridge University Press Cambridge.
- [8] DeVille, R. L., Peskin, C. S., 2008. Synchrony and asynchrony in a fully stochastic neural network. *Bulletin of mathematical biology* 70 (6), 1608–1633.
- [9] Dijkstra, E. W., 1982. Self-stabilization in spite of distributed control. In: *Selected writings on computing: a personal perspective*. Springer, pp. 41–46.
- [10] Dolev, S., 2000. *Self-stabilization*. MIT press.
- [11] Enright, J. T., 1980. Temporal precision in circadian systems: a reliable neuronal clock from unreliable components? *Science* 209 (4464), 1542–1545.
- [12] Ermentrout, G., Kopell, N., 1990. Oscillator death in systems of coupled neural oscillators. *SIAM Journal on Applied Mathematics* 50 (1), 125–146.
- [13] Fisch, R., 1990. Cyclic cellular automata and related processes. *Physica D: Nonlinear Phenomena* 45 (1), 19–25.
- [14] Fisch, R., Gravner, J., Griffeath, D., 1993. Metastability in the greenberg-hastings model. *The Annals of Applied Probability*, 935–967.

- [15] Glass, L., Mackey, M. C., 1979. A simple model for phase locking of biological oscillators. *Journal of Mathematical Biology* 7 (4), 339–352.
- [16] Haefner, J. W., 2005. *Modeling Biological Systems:: Principles and Applications*. Springer.
- [17] Herman, T., Ghosh, S., 1995. Stabilizing phase-clocks. *Information Processing Letters* 54 (5), 259–265.
- [18] Klinglmayr, J., Kirst, C., Bettstetter, C., Timme, M., 2012. Guaranteeing global synchronization in networks with stochastic interactions. *New Journal of Physics* 14 (7), 073031.
- [19] Langton, C. G., 1990. Computation at the edge of chaos: phase transitions and emergent computation. *Physica D: Nonlinear Phenomena* 42 (1), 12–37.
- [20] Mesbahi, M., Egerstedt, M., 2010. *Graph theoretic methods in multiagent networks*. Princeton University Press.
- [21] Mirollo, R. E., Strogatz, S. H., 1990. Synchronization of pulse-coupled biological oscillators. *SIAM Journal on Applied Mathematics* 50 (6), 1645–1662.
- [22] Nair, S., Leonard, N. E., 2007. Stable synchronization of rigid body networks. *Networks and Heterogeneous Media* 2 (4), 597.
- [23] Nishimura, J., Friedman, E. J., 2011. Robust convergence in pulse-coupled oscillators with delays. *Physical review letters* 106 (19), 194101.
- [24] Pavlidis, T., 2012. *Biological oscillators: their mathematical analysis*. Elsevier.
- [25] Peskin, C. S., 1975. *Mathematical aspects of heart physiology*. Courant Institute of Mathematical Sciences, New York University New York.
- [26] Reynolds, C. W., 1987. Flocks, herds and schools: A distributed behavioral model. *ACM SIGGRAPH Computer Graphics* 21 (4), 25–34.
- [27] Strogatz, S. H., 2000. From kuramoto to crawford: exploring the onset of synchronization in populations of coupled oscillators. *Physica D: Nonlinear Phenomena* 143 (1), 1–20.
- [28] Strogatz, S. H., 2001. Exploring complex networks. *Nature* 410 (6825), 268–276.
- [29] Sundararaman, B., Buy, U., Kshemkalyani, A. D., 2005. Clock synchronization for wireless sensor networks: a survey. *Ad Hoc Networks* 3 (3), 281–323.
- [30] Tanner, H. G., Jadbabaie, A., Pappas, G. J., 2003. Stable flocking of mobile agents, part i: Fixed topology. In: *Decision and Control, 2003. Proceedings. 42nd IEEE Conference on*. Vol. 2. IEEE, pp. 2010–2015.
- [31] Tatenos, T., Robinson, H., 2007. Phase resetting curves and oscillatory stability in interneurons of rat somatosensory cortex. *Biophysical Journal* 92 (2), 683–695.
- [32] Winfree, A. T., 1967. The geometry of biological time. Vol. 12.
- [33] Winfree, A. T., 1967. Biological rhythms and the behavior of populations of coupled oscillators. *Journal of theoretical biology* 16 (1), 15–42.
- [34] Wolf-Gladrow, D. A., 2000. *Lattice-gas cellular automata and lattice Boltzmann models: An Introduction*. No. 1725. Springer.
- [35] Wolfram, S., 1984. Cellular automata as models of complexity. *Nature* 311 (5985), 419–424.